CLAIMS

What is claimed is:

1. A method comprising:

receiving a data segment representing a digitized segment of voice data associated with a voice channel;

assigning a unique segment ID to the voice channel associated with the digitized segment of voice data; and

arranging a set of four segment IDs and a corresponding set of four data segments into a quad segment so that the four segment IDs and each of the four data segments are explicitly aligned on an eight-byte boundary.

2. The method of claim 1, further comprising:

prepending a local area network (LAN) header to the quad segment to create a multi-channel voice packet; and

transmitting the multi-channel voice packet over a local area network (LAN).

3. The method of claim 2, wherein the LAN is an Ethernet, and the LAN header is a media access control (MAC) header.

4. The method of claim 2, wherein the LAN is an InfiniBand ® system network.

5. The method of claim 2, further comprising aggregating as many quad segments into the multi-channel voice packet as possible so that a size of the multi-channel voice packet does not exceed the maximum size for the LAN.

6. The method of claim 1, wherein the digitized segment of voice data is at least one sample of pulse-code modulated (PCM) voice data.

7. The method of claim 6, wherein the at least one sample of PCM voice data is one byte in length and represents substantially 125 microseconds of voice data.

8. The method of claim 6, wherein the digitized segment of voice data comprises eight samples of PCM voice data for a total of eight bytes in length.

9. The method of claim 1, wherein the data segment is obtained from a time-division multiplexed (TDM) stream of voice data.

10. The method of claim 1, wherein the data segment is obtained from a asynchronous transfer multiplexed (ATM) stream of voice data.

11. The method of claim 1, wherein the data segment is obtained from a payload of a data packet.

12. The method of claim 11, wherein the data packet is an Real Time Protocol (RTP) packet.

13. The method of claim 1, wherein the unique segment ID is 2 bytes in length, the unique segment ID explicitly identifying the voice channel without reference to other data.

14. A computer-readable medium having executable instructions to cause a computer to perform a method comprising:

receiving a data segment representing a digitized segment of voice data associated with a voice channel;

assigning a unique segment ID to the voice channel associated with the digitized segment of voice data; and

arranging a set of four segment IDs and a corresponding set of four data segments into a quad segment so that the four segment IDs and each of the four data segments are explicitly aligned on an eight-byte boundary.

15. The computer-readable medium of claim 14, wherein the method further comprises:

prepending a local area network (LAN) header to the quad segment to create a multi-channel voice packet; and

transmitting the multi-channel voice packet over a local area network (LAN).

16. The computer-readable medium of claim 15, wherein the LAN is an Ethernet, and the LAN header is a media access control (MAC) header.

17. The computer-readable medium of claim 15, wherein the LAN is an InfiniBand ® system network.

18. The computer-readable medium of claim 15, wherein the method further comprises aggregating as many quad segments into the multi-channel voice packet as possible so that a size of the multi-channel voice packet does not exceed the maximum size for the LAN.

19. The computer-readable medium of claim 14, wherein the digitized segment of voice data is at least one sample of pulse-code modulated (PCM) voice data.

20. The computer-readable medium of claim 19, wherein the at least one sample of PCM voice data is one byte in length and represents substantially 125 microseconds of voice data.

21. The computer-readable medium of claim 20, wherein the digitized segment of voice data comprises eight samples of PCM voice data for a total of eight bytes in length.

22. The computer-readable medium of claim 14, wherein the data segment is obtained from a time-division multiplexed (TDM) stream of voice data.

23. The computer-readable medium of claim 14, wherein the data segment is obtained from a asynchronous transfer multiplexed (ATM) stream of voice data.

24. The computer-readable medium of claim 14, wherein the data segment is obtained from a payload of a data packet.

25. The computer-readable medium of claim 24, wherein the data packet is an Real Time Protocol (RTP) packet.

26. The computer-readable medium of claim 14, wherein the unique segment ID is 2 bytes in length, the unique segment ID explicitly identifying the voice channel without reference to other data.

27. An apparatus comprising:

a data segment receiver to receive a data segment having a digitized segment of voice data;

a voice channel identifier to determine which of a plurality of voice channels is associated with the data segment and to generate an associated segment ID;

a data segment aggregator to arrange a set of four contiguous segment IDs followed by a corresponding set of four contiguous data segments into a quad segment so that the set of four contiguous segment IDs and each of the four contiguous data segments of the quad segment are explicitly aligned on an eight-byte boundary.

28. The apparatus of claim 27, further comprising:

a multi-channel voice packet generator to prepend a local area network (LAN) header to the quad segment to create a multi-channel voice packet; and

a packet transmitter to transmit the multi-channel voice packet over a local area network (LAN).

29. The apparatus of claim 28, wherein the LAN is an Ethernet, and the LAN header is a media access control (MAC) header.

30. The apparatus of claim 28, wherein the LAN is an InfiniBand ® system network.

31. The apparatus of claim 28, wherein the multi-channel voice packet generator further aggregates as many of the quad segments into the multi-channel voice packet as possible so that a size of the multi-channel voice packet does not exceed the maximum packet size for the LAN.

32. The apparatus of claim 28, wherein the digitized segment of voice data is at least one sample of pulse-code modulated (PCM) voice data.

33. The apparatus of claim 32, wherein the at least one sample of PCM voice data is one byte in length and represents substantially 125 microseconds of voice data.

34. The apparatus of claim 32, wherein the digitized segment of voice data comprises eight samples of PCM voice data for a total of eight bytes in length.

35. The apparatus of claim 27, wherein the data segment is obtained from a time-division multiplexed (TDM) stream of voice data.

36. The apparatus of claim 27, wherein the data segment is obtained from an asynchronous transfer multiplexed (ATM) stream of voice data.

37. The apparatus of claim 27, wherein the data segment is obtained from a payload of a data packet.

38. The apparatus of claim 37, wherein the data packet is a Real Time Protocol (RTP) packet.

39. The apparatus of claim 27, wherein the segment ID is 2 bytes in length, the segment ID explicitly identifying the voice channel without reference to other data.

40. A computer-readable medium having stored thereon a data structure, the data structure comprising:

a segment ID representing an identification of a voice channel; and

a data segment representing a digitized segment of voice data associated with the voice channel, wherein the segment ID and the data segment are each positioned to align on an 8-byte boundary.

41. The computer-readable medium of claim 40, wherein four consecutive segment IDs are followed by four corresponding consecutive data segments to form a quad segment, wherein the quad segment is positioned so that the four segment IDs together align on an 8-byte boundary and each of the corresponding four data segments align on an 8-byte boundary.

42. The computer-readable medium of claim 41, wherein the data structure further comprises a local area network (LAN) header representing a destination address associated with the voice channel.

43. The computer-readable medium of claim 41, wherein the segment ID is 2 bytes in length, the segment ID explicitly identifying the voice channel without reference to other data.

44. The computer-readable medium of claim 42, wherein the data structure comprises as many of the quad segments as possible without exceeding the maximum length allowed for the LAN associated with the LAN header.

45. A method comprising:

means for receiving a data segment representing a digitized segment of voice data;

means for assigning a unique segment ID to the voice channel associated with the digitized segment of voice data; and

means for arranging a set of four segment IDs and a corresponding set of four data segments into a quad segment so that the four segment IDs and each of the four data segments are explicitly aligned on an eight-byte boundary.

46. The method of claim 45, further comprising:

means for prepending a local area network (LAN) header to the quad segment to create a multi-channel voice packet; and

means for transmitting the multi-channel voice packet over a local area network (LAN).

47. The method of claim 46, wherein the LAN is an Ethernet, and the LAN header is a media access control (MAC) header.

48. The method of claim 47, further comprising means for aggregating as many of the quad segments into the multi-channel voice packet as possible so that a size of the multi-channel voice packet does not exceed the maximum packet size on the LAN.